

An Efficient Implementation of On-Board Cloud Detection on a Reconfigurable Computer

Esam El-Araby

*The George
Washington University
Washington, DC 20037
USA*

esam@gwu.edu

Mohamed Taher

*The George
Washington University
Washington, DC 20037
USA*

mtaher@gwu.edu

Tarek El-Ghazawi

*The George
Washington University
Washington, DC 20037
USA*

tarek@gwu.edu

Jacqueline Le Moigne

*NASA Goddard Space Flight
Center (GSFC)
Greenbelt, MD 20771
USA*

lemoine@backserv.gsfc.nasa.gov

Abstract – The presence of cloud contamination can hinder the use of satellite data, and this requires a cloud detection process to mask out cloudy pixels from further processing. The trend for remote sensing satellite missions has always been towards smaller size, lower cost, more flexibility, and higher computational power. Reconfigurable Computers (RCs) combine the flexibility of traditional microprocessors with the power of Field Programmable Gate Arrays (FPGAs). Therefore, RCs are a promising candidate for on-board preprocessing.

This paper presents the design and implementation of an RC-based real-time cloud detection system. We investigate the potential of using RCs for on-board preprocessing by prototyping the Landsat 7 ETM+ ACCA algorithm on one of the state-of-the art reconfigurable platforms, SRC-6E.

Although a reasonable amount of investigations of the ACCA cloud detection algorithm using FPGAs has been reported in the literature, very few details/results were provided and/or limited contributions were accomplished. Our work has been proven to provide higher performance and higher detection accuracy.

I. INTRODUCTION

The trend for remote sensing satellite missions has always been towards smaller size, lower cost, and more flexibility. On-board processing, as a solution, permits a good utilization of expensive resources. Instead of storing and forwarding all captured images, data processing can be performed on-orbit prior to downlink resulting in the reduction of communication bandwidth as well as in simpler and faster subsequent computations to be performed on ground stations. Consequently, on-board processing can reduce the cost and the complexity of the On-The-Ground/Earth processing systems. Furthermore, it enables autonomous decisions to be taken on-board which can potentially reduce the delay between image capture, analysis and action. This leads to faster critical decisions which are crucial for future reconfigurable web sensors missions as well as planetary exploration missions.

The presence of cloud contamination can hinder the use of satellite data, and this requires a cloud detection

process to mask out cloudy pixels from further processing. The Landsat 7 ETM+ (Enhanced Thematic Mapper) ACCA (Automatic Cloud Cover Assessment) algorithm [1] is a compromise between the simplicity of earlier Landsat algorithms, e.g. ACCA for Landsat 4 and 5, and the complexity of later approaches such as the MODIS (Moderate Resolution Imaging Spectroradiometer) cloud mask.

Reconfigurable Computers (RCs) combine the flexibility of traditional microprocessors with the power of Field Programmable Gate Arrays (FPGAs). These platforms have always been reported to outperform the conventional platforms in terms of throughput and processing power within the domain of cryptography, and image processing applications [2]. In addition, they are characterized by lower form/wrap factors compared to parallel platforms, and higher flexibility than ASIC solutions. Therefore, RCs are a promising candidate for on-board preprocessing. The SRC-6E Reconfigurable Computer is one example of this category of hybrid computers [3] and is used here for this purpose.

This paper presents the design and implementation of an RC-based real-time cloud detection system. We investigate the potential of using RCs for on-board preprocessing by prototyping the Landsat 7 ETM+ ACCA algorithm on one of the state-of-the art reconfigurable platforms, SRC-6E.

Although a reasonable amount of investigations of the ACCA cloud detection algorithm using FPGAs has been reported in the literature [4], [5], very few details/results were provided and/or limited contributions were accomplished. Our work is unique in providing higher performance and higher detection accuracy.

II. THE AUTOMATIC CLOUD COVER ASSESSMENT (ACCA)

Theory of Landsat 7 ETM+ ACCA algorithm is based on the observation that clouds are highly reflective and cold. The high reflectivity can be detected in the visible, near- and mid- IR bands. The thermal properties of clouds can be detected in thermal

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 25 JUL 2005		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE An Efficient Implementation of On-Board Cloud Detection on a Reconfigurable Computer				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The George Washington University Washington, DC 20037 USA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM001850, 2005 IEEE International Geoscience and Remote Sensing Symposium Proceedings (25th) (IGARSS 2005) Held in Seoul, Korea on 25-29 July 2005. , The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 4	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

IR band. Table I presents the bands wavelengths and its detection features.

TABLE I
LANDSAT 7 ETM+ BANDS

Band	Wavelength (μm)	Detection Features
2 (green)	0.525 - 0.605	- Measures green reflectance - Vegetation discrimination
3 (red)	0.630 - 0.690	- Measures Chlorophyll absorption - Plant Species differentiation
4 (near IR)	0.775 - 0.900	- Determines soil moisture level - Delineating water bodies and distinguishing vegetation types
5 (mid IR)	1.55 - 1.75	- Supplies information about vegetation and soil moisture - Differentiation of snow from clouds
6 (thermal IR)	10.4 - 12.5	- Thermal mapping to Brightness Temperatures

The Landsat 7 ETM+ ACCA algorithm recognizes clouds by analyzing the scene twice. In the first pass six filters are utilized for this purpose, see Table II. Omission errors are expected. The goal of pass-one is to develop a reliable cloud signature for use in pass-two where the remaining clouds are identified. Commission errors, however, create algorithm failure and must be minimized. Three categories result from pass-one: clouds, non-clouds, and an ambiguous group that is revisited in pass-two.

In pass-two of the algorithm, descriptive statistics are calculated from band 6 to describe the cloud category: these include mean temperature, standard deviation and distribution skew. New band 6 thresholds are developed from these statistics. Only the thermal band is examined during pass-two in order to capture the remaining clouds. Image pixels that fall below the new threshold qualify as cloud pixels. After pass-two processing, cloud cover results from both pass-one and pass-two are compared. Extreme differences are indicative of cloud signature corruption. When this occurs, pass-two results are ignored and all results are taken from pass-one.

TABLE II
PASS-ONE FILTERS

Filter	Function
Band 3 threshold	Eliminates dark images
(Band2 - Band5)/ (Band2 + Band5)	Eliminates many types of snow
Band 6 Threshold	Eliminates warm image features
(1 - Band5) * Band6	Eliminates numerous categories including ice
Band 4/3 ratio	Eliminates bright vegetation and soil
Band 4/5 ratio	Eliminates rocks and desert

During processing, a cloud mask is created. After the two ACCA passes, a filter is applied to the cloud mask to fill in cloud holes. This filtering operation works by examining each non-cloud pixel in the mask. If 5 out of the 8 neighbors are clouds then the pixel is

reclassified as cloud. The final cloud cover percentage for the image is calculated based on the filtered cloud mask.

III. SRC-6E RECONFIGURABLE COMPUTER

A. Hardware Architecture

SRC-6E platform consists of two general-purpose microprocessor boards and one MAP[®] reconfigurable processor board. Each microprocessor board is based on two 1 GHz Pentium 3 microprocessors. The SRC MAP board consists of two MAP reconfigurable processors.

Overall, the SRC-6E system provides a 1:1 microprocessor to FPGA ratio. Microprocessor boards are connected to the MAP board through the SNAP[®] interconnect. SNAP card plugs into the DIMM slot on the microprocessor motherboard [3].

Hardware architecture of the SRC MAP processor is shown in Fig. 1. This processor consists of two programmable User FPGAs, six 4 MB banks of the on-board memory (OBM), and a single Control FPGA. The FPGAs are all Xilinx Virtex II-6000-4.

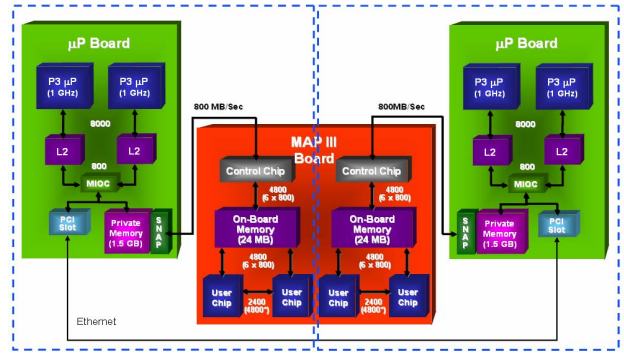


Fig. 1 Hardware Architecture of SRC-6E

B. Programming Model

The SRC-6E has a similar compilation process as a conventional microprocessor-based computing system, but needs to support additional tasks in order to produce

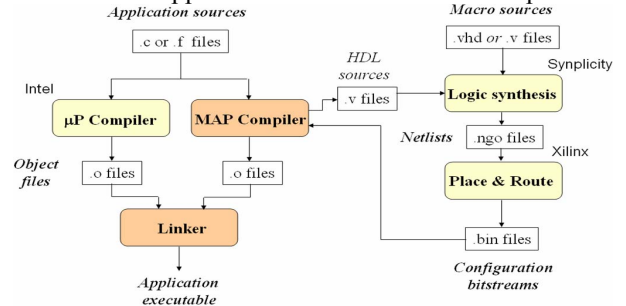


Fig. 2 SRC Compilation Process

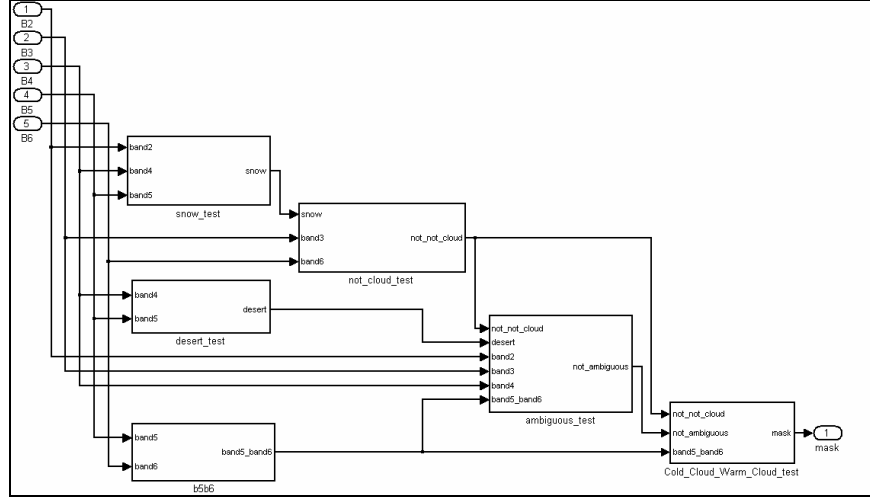


Fig. 3 Pass-One Module

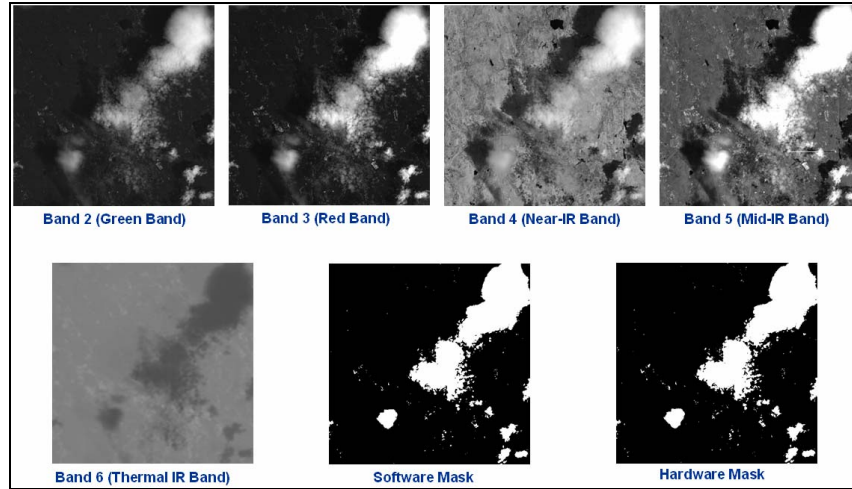


Fig. 4 Detection Accuracy

logic for the MAP reconfigurable processor, as shown in Fig. 2. Since users often wish to extend the built-in set of operators, the compiler allows users to integrate their own VHDL/Verilog macros.

IV. ALGORITHM IMPLEMENTATION AND EXPERIMENTAL RESULTS

The ACCA algorithm adapted for Landsat 7 ETM+ data has been implemented both in C and Matlab, and pass-one has been implemented and synthesized for the Xilinx XC2V6000 FPGA on SRC-6E.

Fig 3. shows the implementation of pass-one of the algorithm. This module reads the five band-image and feeds them to six filters generating the output mask. The function of each mask is listed in Table II.

Our goal for the implementation was to achieve a high performance and high accuracy as compared to what has been reported in [4] and [5]. Therefore, the


only constraint to our design was the processing speed, as measured by throughput. This constraint is approached through full-pipelining of the design.

Many of the tests in pass-one are threshold tests of ratio values, such as the snow test. We found out that it was more efficient, in terms of the required resources, to multiply one value by the threshold, and compare with the other value, instead of performing the division then comparing against the threshold.


The design was developed in VHDL, synthesized, placed and routed, and was found to occupy approximately 7% of the available logic resources on the chip. This enabled us to instantiate eight concurrent processing engines of the design in the same chip, which increased the performance to eight folds of what we expected. The total resources utilization for the eight engine version was approximately 57%, which leaves plenty of room for more engines and/or processing functions to be implemented on the same

chip. The maximum operational clock speed of the design is 100MHz which resulted in 4000 Megapixels/sec (5 inputs x 8 engines x 100MHz) as data input/consumption rate. Furthermore, the data output/production rate was 800 Megapixels/sec (1 output x 8 engines x 100MHz).

Fig. 4 shows the results obtained from a 2.8GHz Intel Xeon processor and from SRC-6E. The hardware (fixed point) implementation provided a very high performance (28 times faster) compared to the 2.8GHz Xeon implementation. We have obtained approximately an ideal detection accuracy (0% detection error) as compared to software floating point reference results. These results were achieved by performing the internal algorithm filtering with full-precision (23-bit) fixed point arithmetic.




(a) Throughput




(b) Speedup

Fig. 5 Hardware-to-Hardware Performance



(a) Execution Time



(b) Speedup

Fig. 6 Hardware-to-Software Performance

Our results have been compared to previous work by Williams et al. [5]. Fig. 5 shows those comparisons. Three hardware implementations have been compared: 1X (one engine is instantiated), 2X (2 engines are instantiated), and 8X (8 engines are instantiated). Our implementations achieve a speedup of up to 16 times compared to those reported in [5].

The superiority of RCs over traditional platforms for cloud detection is demonstrated through the performance plots shown in Fig. 6. Our implementations achieve a speedup of up to 28.3 compared to 2.8GHz Intel Xeon processor.

V. CONCLUSIONS

This paper presents the design and implementation of an RC-based real-time cloud detection system. We investigated the potential of using RCs for on-board preprocessing by prototyping the Landsat 7 ETM+ ACCA algorithm on one of the state-of-the art reconfigurable platforms, SRC-6E.

Our work has been proven to provide higher performance and higher detection accuracy than previously reported results. The higher performance is achieved through full-pipelining and superscaling (up to 8 concurrent engines), and thus achieving 4000 Megapixels/sec as a data consumption rate and 800 Megapixel/sec as a data production rate. In addition, the performance has been compared to similar hardware implementation and proved to achieve as high as 16 folds speedup. The speedup compared to a 2.8GHz Xeon implementation has been 28 folds higher. On the other hand, the detection accuracy has been verified against software floating-point reference implementation, and the results revealed identical results.

REFERENCES

- [1] R.R. Irish, "Landsat 7 Automatic Cloud Cover Assessment," Algorithms for Multispectral, Hyperspectral and Ultraspectral Imagery VI, SPIE, Orlando, FL., USA, 24-26 April 2000, pp.348-355.
- [2] E. El-Araby, T. El-Ghazawi, J. Le Moigne, and K. Gaj, "Wavelet Spectral Dimension Reduction of Hyperspectral Imagery on a Reconfigurable Computer," IEEE International Conference on Field-Programmable Technology, FPT 2004, Brisbane, Australia, December 2004.
- [3] "SRC-6E C-Programming Environment Guide", SRC Computers, Inc. 2004.
- [4] R.S. Basso, J. Le Moigne, S. Veuella, and R.R. Irish, "FPGA Implementation for On-Board Cloud Detection," International Geoscience and Remote Sensing Symposium. Hawaii, 20-24 July 2000.
- [5] J.A. Williams, A.S. Dawood, S.J. Visser, "FPGA-based Cloud Detection for Real-Time Onboard Remote Sensing," Proceedings of IEEE International Conference on Field-Programmable Technology (FPT 2002), 16-18 Dec. 2002, pp.110 – 116.